

Cloudera AI 1.5.4

Release Notes

Date published: 2020-07-16

Date modified: 2024-12-16

CLOUDERA

<https://docs.cloudera.com/>

Legal Notice

© Cloudera Inc. 2025. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

Contents

What's New	4
Known Issues	5
Fixed Issues	8
Cumulative Hotfixes: Cloudera AI on premises	8
Cloudera AI on premises 1.5.4 CHF1.....	9
Fixed issues in 1.5.4 CHF1.....	9
Known issues in 1.5.4 CHF1.....	10
Cloudera AI on premises 1.5.4 CHF2.....	10
Cloudera AI on premises 1.5.4 CHF3.....	10
Fixed issues in 1.5.4 CHF3.....	10
Cloudera AI on premises 1.5.4 SP2 CHF1.....	11
Fixed issues in 1.5.4 SP2 CHF1.....	11
Repository Locations for 1.5.4 SP2 CHF1.....	12
Service Packs: Cloudera AI on premises	12
Cloudera AI on premises 1.5.4 SP1.....	12
What's new in 1.5.4 SP1.....	12
Known issues in 1.5.4 SP1.....	12
Fixed issues in 1.5.4 SP1.....	19
Cloudera AI on premises 1.5.4 SP2.....	19
What's new in 1.5.4 SP2.....	19
Known issues in 1.5.4 SP2.....	19
Fixed issues in 1.5.4 SP2.....	20

What's New

Cloudera on premises 1.5.4 includes the following features for Cloudera AI.

New features

Cloudera AI Service Accounts are available in Cloudera AI on premises

In Cloudera AI, the Kerberos principal for the Service Account may not be the same as your login information. Therefore, ensure you provide the Kerberos identity when you sign in to the Service Account. For more information, see [Authenticating Hadoop for Cloudera AI Service Accounts](#).

Cloudera AI Registry is available in Cloudera on premises

Cloudera AI Registry is now generally available (GA) in on premises. Cloudera AI Registry in on premises uses Apache Ozone to store model artifacts. For creating a Cloudera AI Registry you need the Ozone S3 gateway endpoint, the Ozone access key, and the Ozone secret key.

If you deploy Cloudera AI Registry in an environment that contains one or more Cloudera AI Workbenches, you must synchronize the Cloudera AI Registry with the workbenches. For more information, see [Prerequisites for creating Cloudera AI Registry](#) and [Synchronizing the Cloudera AI Registry with a workbench](#).

Heterogeneous GPU usage

When using heterogeneous GPU clusters to run sessions and jobs, the available GPU accelerator labels need to be selected during workload creation. For more information, see [Heterogeneous GPU clusters](#).

Data connections without auto discovery

Cloudera AI is a flexible, open platform, supporting connections to many data sources. The provided code samples demonstrate how to access local data for Cloudera AI workloads. For more information, see [Connecting to Cloudera Data Warehouse](#).

Spark Log4j Configuration

Cloudera AI allows you to update Spark's internal logging configuration on a per-project basis. Spark logging properties can be customized for every session, and job with a default file path found at the root of your project. You can also specify a custom location with a custom environment variable. For more information, see [Spark Log4j Configuration](#).

ML Metrics Collector service

The Metrics Collector service gathers data about how users and groups use resource quota, like how much CPU, Memory and GPU capacity (if any) is allocated, and what the users or groups utilize from that. The Metrics Collector service is running by default, but to collect data about resource quota metrics, you need to enable the Quota Management feature. For more information, see [ML Metrics Collector Service overview](#).

Quota Management for group level

Quota Management Technical Preview (TP) release enables you to control how resources are allocated within your Cloudera AI Workbench on user and on group level. Yunikorn Gang Scheduling is also available, which is the default scheduling mechanism in Cloudera AI. For more information, see [Quota Management overview](#) and [Yunikorn Gang Scheduling](#).

Restarting a failed AMP setup

You can now retry failed AMP deployment steps and continue the AMP setup to handle intermittent and configuration issues. For more information, see [Restarting a failed AMP setup](#).

New Hadoop CLI Runtime Addon versions are available

The HadoopCLI 7.1.8.3-601 Runtime Addon is released for the Cloudera on premises.

Related Information

[Known Issues and Limitations](#)

Known Issues

You might run into some known issues while using Cloudera AI on premises.

Using dollar character in environment variables in Cloudera AI

Environment variables with the dollar (\$) character are not parsed correctly by Cloudera AI. For example, if you set `PASSWORD="pass$123"` in the project environment variables, and then try to read it using the `echo` command, the following output will be displayed: `pass23`

Workaround: Use one of the following commands to print the \$ sign:

```
echo 24 | xxd -r -p
or
echo JAo= | base64 -d
```

Insert the value of the environment variable by wrapping it in the command substitution using `$()` or ```. For example, if you want to set the environment variable to `ABC$123`, specify:

```
ABC$(echo 24 | xxd -r -p)123
or
ABC`echo 24 | xxd -r -p`123
```

DSE-37827: Jupyter's RTC extension throws an error and notebooks become unusable

In certain cases, Jupyter's RTC (Real Time Collaboration) extension may cause errors claiming either that other sessions are active, or that other processes have accessed the notebook files. After these errors, the notebook becomes unusable due to the error messages and the Cloudera AI session needs to be restarted.

Workaround:

You must disable the Jupyter RTC extension by performing the following tasks:

1. Create a Session.
2. Open the terminal.
3. Enter `nano /home/cdsw/.jupyter/labconfig/page_config.json`.
4. Add the following lines to the file:

```
{
  "disabledExtensions": {
    "@jupyter/collaboration-extension": true
  },
  "lockedExtensions": {
    "@jupyter/collaboration-extension": true
  }
}
```

5. Save and close the file.

DSE-36718: Disable auto synchronization feature for users and teams

The automated team and user synchronization feature is disabled. Newly installed or upgraded workbenches do not have the automatic synchronization option in the Cloudera AI UI.

Workaround: none

DSE-36759: AMPs and Feature Announcement sections do not work in NTP setups

Cloudera AI on premises setups with Non Transparent Proxy do not function properly, that affects Cloudera Accelerators for Machine Learning Projects and Feature Announcements. The home page freezes, the feature announcement displays error message, and the AMPs do not load.

Workaround:

To avoid the home page freeze copy the following environment variables from the web deployment, and add them to the environment section of the API deployments:

- HTTP_PROXY
- HTTPS_PROXY
- NO_PROXY
- http_proxy
- https_proxy
- no_proxy



Note: The values are the same for both the lower and upper case variants.

DSE-32943: Enabling Service Accounts

Teams in the Cloudera AI Workbench can only run workloads within team projects with the Run as option for service accounts if they have previously manually added service accounts as a collaborator to the team.

DSE-35013: First Cloudera AI Workbench creation fails

On RHEL 8.8, during the first Cloudera AI Workbench installation on GPU with Cloudera Embedded Container Service external registry, pods might get stuck in the init or CrashLoop state.

First-time workbench installation is expected to fail. Consider this as a test workbench, and apply the following manual workaround for creating subsequent workbenches:

1. Restart or delete the pods which are in init or CrashLoop state in the test workbench.
2. Once all pods are in the running state, create new workbenches as needed.
3. Delete the test workbench from the Cloudera AI UI if no longer needed.

OPX-4603: Buildkit in Cloudera Embedded Container Service in Cloudera AI on premises

Issue: BuildKit was introduced in Cloudera Embedded Container Service for building images of models and experiments. BuildKit is a replacement for Docker, which was previously used to build images of Cloudera AI's models and experiments in Cloudera Embedded Container Service. Buildkit is only for OS RHEL8.x and CentOS 8.x.

Buildkit in Cloudera AI on premises 1.5.2 is a Technical Preview feature. Hence, having Docker installed on the nodes/hosts is still mandatory for models and experiments to work smoothly. Upcoming release will be completely eliminating the dependency of Docker on the nodes.

Workaround: None.

DSE-32285: Migration: Migrated models are failing due to image pull errors

Issue: After CDSW to Cloudera AI migration (on-premises) via full-fledged migration tool, migrated models on Cloudera AI Workbench on premises fails on initial deployment. This is because the initial model deployment tries to pull images from on-premises's registry.

Workaround: Redeploy the migrated model. As this involves the build and deploy process, the image will be built, pushed to the on premises Cloudera AI Workbench's configured registry, and then the same image will be consumed for further usage.

DSE-28768: Spark Pushdown is not working with Scala 2.11 runtime

Issue: Scala and R are not supported for Spark Pushdown.

Workaround: None.

DSE-32304: On Cloudera AI on premises on Cloudera Embedded Container Service terminal and ssh connections can terminate

Issue: In Cloudera on premises on Cloudera Embedded Container Service, Cloudera AI Terminal and SSH connections can terminate after an uncertain amount of time, usually after 4-10 minutes. This issue affects the usage of local IDEs to work with Cloudera AI, as well as any customer application using a websocket connection.

Workaround: None.

DSE- 35251: Web pod crashes if a project forking takes more than 60 minutes

The web pod crashes if a project forking takes more than 60 minutes. This is because the timeout is set to 60 minutes using the `grpc_git_clone_timeout_minutes` property. The following error is displayed after the web pod crash:

```
2024-04-23 22:52:36.384 1737 ERROR AppServer.VFS.grpc
crossCopy grpc error data = [{"error": "1"},
{"code": 4, "details": "2", "metadata": "3"}, "Deadline exceeded", {}]
["Error: 4 DEADLINE_EXCEEDED: Deadline exceeded\n
at callErrorFromStatus (/home/cdswint/services/web/node_modules/
@grpc/grpc-js/build/src/call.js:31:19)\n at Object.onReceiveS
tatus (/home/cdswint/services/web/node_modules/@grpc/grpc-js/bui
ld/src/client.js:192:76)\n at Object.onReceiveStatus (/home/c
dswint/services/web/node_modules/@grpc/grpc-js/build/src/client-
interceptors.js:360:141)\n at Object.onReceiveStatus (/home/c
dswint/services/web/node_modules/@grpc/grpc-js/build/src/client-
interceptors.js:323:181)\n at /home/cdswint/services/web/node
_modules/@grpc/grpc-js/build/src/resolving-call.js:94:78\n at
process.processTicksAndRejections (node:internal/process/task_q
ueues:77:11)\nfor call at\n at ServiceClientImpl.makeUnaryReq
uest (/home/cdswint/services/web/node_modules/@grpc/grpc-js/buil
d/src/client.js:160:34)\n at ServiceClientImpl.crossCopy (/ho
me/cdswint/services/web/node_modules/@grpc/grpc-js/build/src/mak
e-client.js:105:19)\n at /home/cdswint/services/web/server-di
st/grpc/vfs-client.js:235:19\n at new Promise (<anonymous>)\n
at Object.crossCopy (/home/cdswint/services/web/server-dist/
grpc/vfs-client.js:234:12)\n at Object.crossCopy (/home/cdswi
nt/services/web/server-dist/models/vfs.js:280:38)\n at projec
tForkAsyncWrapper (/home/cdswint/services/web/server-dist/models/
projects/projects-create.js:229:19)"]
node:internal/process/promises:288
triggerUncaughtException(err, true /* fromPromise */);
^Error: 4 DEADLINE_EXCEEDED: Deadline exceeded
at callErrorFromStatus (/home/cdswint/services/web/n
ode_modules/@grpc/grpc-js/build/src/call.js:31:19)
at Object.onReceiveStatus (/home/cdswint/services/web/n
ode_modules/@grpc/grpc-js/build/src/client.js:192:76)
at Object.onReceiveStatus (/home/cdswint/services/we
b/node_modules/@grpc/grpc-js/build/src/client-interceptors.js:36
0:141)
at Object.onReceiveStatus (/home/cdswint/services/web/n
ode_modules/@grpc/grpc-js/build/src/client-interceptors.js:323:1
81)
at /home/cdswint/services/web/node_modules/@grpc/grpc-
js/build/src/resolving-call.js:94:78
at process.processTicksAndRejections (node:internal/p
rocess/task_queues:77:11)
for call at
at ServiceClientImpl.makeUnaryRequest (/home/cdswint/s
ervices/web/node_modules/@grpc/grpc-js/build/src/client.js:160:3
4)
```

```

    at ServiceClientImpl.crossCopy (/home/cdswint/services/
web/node_modules/@grpc/grpc-js/build/src/make-client.js:105:19)
    at /home/cdswint/services/web/server-dist/grpc/vfs-cli
ent.js:235:19
    at new Promise (<anonymous>)
    at Object.crossCopy (/home/cdswint/services/web/server-
dist/grpc/vfs-client.js:234:12)
    at Object.crossCopy (/home/cdswint/services/web/server-
dist/models/vfs.js:280:38)
    at projectForkAsyncWrapper (/home/cdswint/services/web/
server-dist/models/projects/projects-create.js:229:19) {
  code: 4,
  details: 'Deadline exceeded',
  metadata: Metadata { internalRepr: Map(0) {}, options:
{} }
}

```

Workaround: Increase the timeout limit, for example, to 120 minutes, using the `grpc_git_clone_timeout_minutes` property.

```

UPDATE site_config SET grpc_git_clone_timeout_minutes = <NEW
VALUE>;

```

Fixed Issues

This section lists issues fixed in this release for Cloudera AI on premises.

DSE-24333: Embedded Spark UI fails to load in Spark pushdown mode for Spark 3

When Spark is run on yarn, the configuration attempts to disable proxying of the Spark UI through the Yarn Resource Manager. This setting was not respected by Spark 3 binaries and resulted in the Could not determine the proxy server for redirection error message when attempting to reach the embedded Spark UI on the driver.

This behavior works again in Spark 3.2.3 with CDH 7.1.7 and Spark 3.3.2 with CDH 7.1.9.

DSE-32887: Runtime updates cannot be enabled from the Cloudera Manager UI Site Administration dashboard

Runtime updates can now be enabled from the Cloudera AI UI Site Administration dashboard.

DSE-29976: Job Arguments fail in PBJ runtimes

PBJ (Powered by Jupyter) Runtimes were not able to access Job arguments as standard command line arguments. To mitigate this, a new environment variable was added to Jobs, called "JOB_ARGUMENTS". This environment variable can be used to access arguments specified in the Job Creation or Job Update UI.

DSE-27285: The calculation of User Resources was incorrect

The **User Resources** tab in the Cloudera AI Workbench UI calculated with the used and available resources for a cluster that also included the resources for the master node. As the workload cannot be scheduled in the master nodes, the master node resources shall not be considered in these calculations. This issue has been solved, the resources for the master are no longer included in the calculations for the User Resources.

Cumulative Hotfixes: Cloudera AI on premises

Review the cumulative hotfixes for Cloudera on premises 1.5.4.

Cloudera AI on premises 1.5.4 CHF1

Review the features, fixes, and known issues in the Cloudera AI 1.5.4 Cumulative hotfix 1 release.

Fixed issues in 1.5.4 CHF1

Review the fixed issues in the Cloudera AI 1.5.4 Cumulative hotfix 1 release.

DSE-37209: Upgrade of Cloudera AI Workbench with external NFS fails

For setups with an external network file system (NFS), the Cloudera AI Workbench upgrade flow was different from setups with internal NFSs. There was an issue with the persistent volume claim (PVC) size in external NFS flows which caused failures during the upgrade path to 1.5.4. Now, the issue has been fixed. For more details, see [Cloudera Customer Advisory 2024-778: Cloudera AI Cloudera on premises Cloudera upgrade to 1.5.4 is failing with external NFS](#) Technical Service Bulletin.

DSE-36078: Spark workload histories are not logged in Spark history server

The Spark history server did not log Spark workload history so far. This issue has been fixed and the Spark workload history is now logged in the base cluster Spark history server.

DSE-37192: Time zone discrepancy for usage API

Some data with respect to the workload run were stored in the local time zone, as a result of which the API returned empty results. By setting all the columns of the dashboard database to UTC time zone, there is no longer time zone discrepancy for the usage API. Note that the data stored prior to the upgrade remains in the previously set time zone.

DSE-36489: Opening existing JupyterLab session redirected to project page

After having logged out from the Cloudera AI UI and having logged back, the opening of an existing JupyterLab session led to the project page instead of opening the JupyterLab UI. This issue has been fixed and after a new login the JupyterLab UI can be opened.

DSE-32204: Improved error messaging during failed Kerberos authentication in session startup

If there was an issue with Kerberos credential authentication, Cloudera AI session startup hung until the regular session timed out and no warning message was displayed for the user. Now, a warning message with additional details on the status of the Kerberos credential authentication is displayed for the user along with a startup failure message after 1 minute.

DSE-36666: The migration readiness check is disrupted when engine readiness check is run

The shell command executor on CDSW could not handle the CDSW to Cloudera AI migration readiness timeout. This issue has been fixed, the timeout has been handled. Additionally, any SQL query-based readiness checks that remain unhandled will be displayed for manual validation.

DSE-36759: AMPs and Feature Announcement sections do not work in NTP setups

Cloudera AI Cloudera on premises setups with non-transparent proxy (NTP) did not function properly, affecting Cloudera Accelerators for Machine Learning Projects and feature announcements. The home page froze, the feature announcement displayed an error message, and the AMPs did not load. This issue has been fixed with the enablement of non-transparent proxy in a Cloudera on premises setup.

DSE-37454: Possible Cloudera AI terminal disconnections

The Cloudera AI terminal tended to disconnect on Cloudera Embedded Container Service because the proxy_read_timeout value on workload pods was set to a default value of 60 seconds. The proxy_read_timeout value has been increased to 121 seconds and the issue of the Cloudera AI terminal disconnecting connections has been resolved.

DSE-37610: Due to version mismatch of the jackson-module-scala_3 the mlops governance pod image failed

Earlier, the version of the jackson-module-scala_3 was updated in the mlopsgovernance/pom.xml file, but it was not updated in the parent/pom.xml file. This issue has been solved and there is no longer version mismatch for the jackson-module-scala_3.

DSE-36748: File upload issue fixed using APIv2

Earlier, when uploading a file to a sub-folder in a project, the file was added to the root folder. This issue has been fixed and now when uploading a file to a subfolder it is added to the project's sub folder.

Known issues in 1.5.4 CHF1

Review the known issues in the Cloudera AI 1.5.4 Cumulative hotfix 1 release.

DSE-36967: Namespace Termination issue when using Portworx storage

There is an issue with Portworx version lower than 3.1.1, as the namespace deletion gets stuck in terminating state. Portworx is not able to cleanly unmount and clean up the underlying resources.

Workaround: The issues is fixed with Portworx version 3.1.1. Upgrade to Portworx version 3.1.1 or to later versions.

DSE-37002: Download API using v2 python swagger client does not work

The download_project_file method does not function as expected. The v2 python swagger only allows content to be downloaded as a string and not in bytes, as described in [Binary string response encoding with Python 3](#).

Workaround: Use download_project_file_v2 as an alternate method.

Cloudera AI on premises 1.5.4 CHF2

There were no features, fixes, and known issues in the Cloudera AI 1.5.4 Cumulative hotfix 2 release.

Cloudera AI on premises 1.5.4 CHF3

Review the features, fixes, and known issues in the Cloudera AI 1.5.4 Cumulative hotfix 3 release.

Fixed issues in 1.5.4 CHF3

Review the fixed issues in the Cloudera AI 1.5.4 Cumulative hotfix 3 release.

DSE-37611: Cloudera AI v2 API deployed application did not inherit user level environment variables and site level environment variables

Applications created using APIv2 did not inherit user level environment variables and site level environment variables. This issue has been solved, and now an application created using APIv2 does not only inherit project level environment variables but also user level environment variables and site level environment variables. This ensures that the behaviour of APIv2 is in synchronisation with APIv1.

DSE-38499: Spark client configurations are not mounted to Spark executor pods

hadoop/conf/ files were mounted only on Spark engines but not to the Spark executors, resulting in a Kerberos issue in Spark applications using third-party file systems like PowerScale (Isilon). This issue has been fixed by including the hadoop_conf_dir in the Spark executor template YAML.

DSE-38627: Modifications for a job's resource profile could not be saved

The issue has been fixed and now the resource profile dropdown allows alternate selections and also the **Job Settings** page persists resource profile selection after saving.

Cloudera AI on premises 1.5.4 SP2 CHF1

Review the features, fixes, and known issues in the Cloudera AI 1.5.4 SP2 Cumulative hotfix 1 release.

Fixed issues in 1.5.4 SP2 CHF1

Review the fixed issues in the Cloudera AI 1.5.4 SP2 Cumulative hotfix 1 release.

DSE-40325: Frequently received 504 timeout error for VFS stat calls

Concurrent project forks have caused 504 timeout errors when navigating the Projects page.

This issue has been resolved.

DSE-41027: Upgrade path from 1.5.0 -> 1.5.2 -> 1.5.4 -> 1.5.4 CHF3 fails with missing key

Due to a strict validation introduced, certain upgrade paths were missing the required key, resulting in upgrade failures. Users upgrading to version 1.5.4 might have encountered this issue if their original Cloudera Data Services on premises version at the time of initial installation was 1.5.0 or 1.5.1.

This issue has been resolved.

DSE-42379: Model registry page and its deletion result in errors on Cloudera AI Registry page

While navigating the Cloudera AI Registry pages, the following error was displayed: Could not interpret page token: provided token did not match parameters.

Additionally, when attempting to delete a model, the model disappeared from the page and an error message appeared: Failed to delete model from model registry &{ } (*models.Error) is not supported by the TextConsumer, can be resolved by supporting TextUnmarshaller interface.

These issues occurred due to failures in page token validation on the Cloudera AI Registry pages.

These issues have been resolved.

DSE-43808: 2.0.44: Cloudera AI Registry workbench pagination issues

While navigating to the **next page** in the Cloudera AI Registry workbench UI, an error was displayed. Additionally, selecting the option to display 25-51 results per page often resulted in showing only 10 items.

These issues have been resolved.

DSE-42480: List of Cloudera AI Registry calls in workbench is invoked with page_size=10 10

When navigating on the Cloudera AI Registry pages to view the models from the workbench, it was always displayed with page_size=10, even if you had previously selected another value.

This issue has been resolved.

DSE-43104: Timezones in Cloudera on premises cause pods to be killed with exit code 34

In an on premises environment, the Kubernetes cluster could be configured to any timezone, leading to timezone discrepancies, particularly for engine pods. This issue affected engine timestamp fields, such as **scheduling_at**, **starting_at**, **running_at**, and **finished_at**, which were reported in varying timezones throughout the infrastructure. As a result, when the base cluster and the nodes were in different timezones, the engine pods were terminated.

This issue has been resolved.

Repository Locations for 1.5.4 SP2 CHF1

The URLs for Cloudera AI on premises 1.5.4 SP2 CHF1 are listed in the following table:

URL Type	Repository Location
Index	<code>https://username:password@archive.cloudera.com/p/cdp-pvc-ds/1.5.4-h19/</code>
Manifest	Repository: <code>https://username:password@archive.cloudera.com/p/cdp-pvc-ds/1.5.4-h19/manifest.json</code>
Parcels	Repository: <code>https://username:password@archive.cloudera.com/p/cdp-pvc-ds/1.5.4-h19/parcels/</code>

Service Packs: Cloudera AI on premises

Review the Service pack features and issues for Cloudera on premises 1.5.4.

Cloudera AI on premises 1.5.4 SP1

Review the features, fixes, and known issues in the Cloudera AI 1.5.4 Service Pack 1 release.

What's new in 1.5.4 SP1

Cloudera on premises 1.5.4 SP1 includes the following features for Cloudera AI.

Exploratory Data Science and Visualization

The Exploratory Data Science and Visualization feature automatically discovers the data sources available for you, from within the standard Cloudera AI user interface. Exploratory Data Science and Visualization enables the Data Scientist to connect to data sources within a project, to explore data, its basic shape and characteristics, to create data sets to be used, to visualize data for understanding its nature, and to create dashboards for sharing purposes.

Known issues in 1.5.4 SP1

You might run into some known issues while using Cloudera AI on premises.

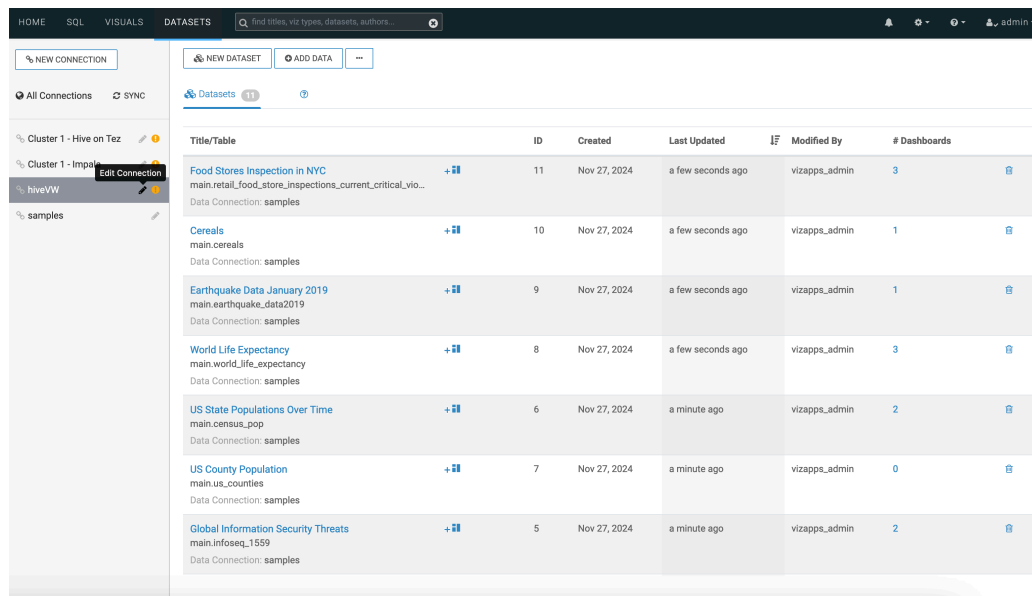
Required workarounds with Cloudera Data Visualization with Hive and Impala

Cloudera Data Visualization on premises for Hive and Impala Virtual Warehouses

For starting Cloudera Data Visualization on premises for Hive and Impala Virtual Warehouses, switch to **LDAP Authentication mode** and add your username and password in the **Edit Data Connection** page:

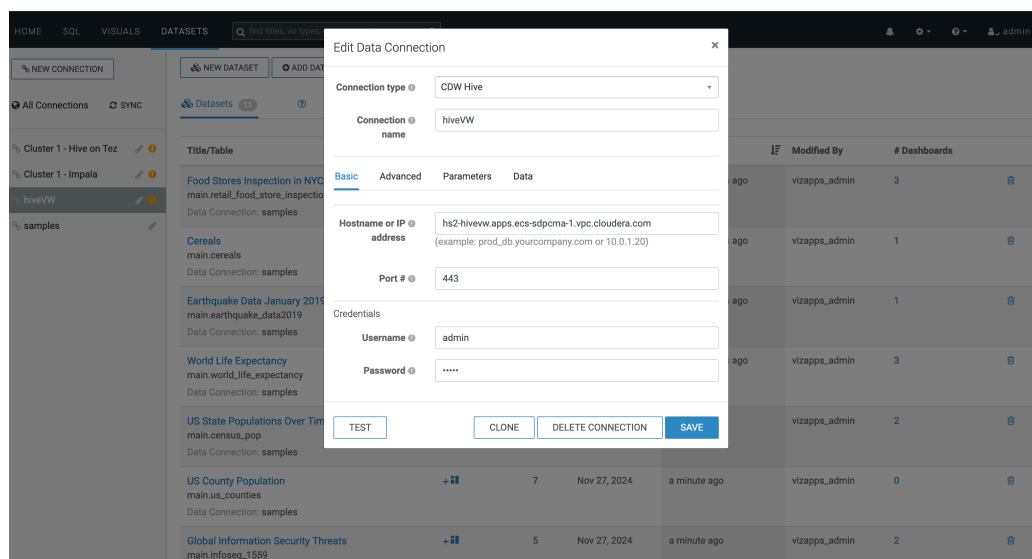
1. Select the required connection in the left navigation and click Edit Connection.

Figure 1: Selecting the required connection for editing



2. Provide your username and password.

Figure 2: Provide username and password

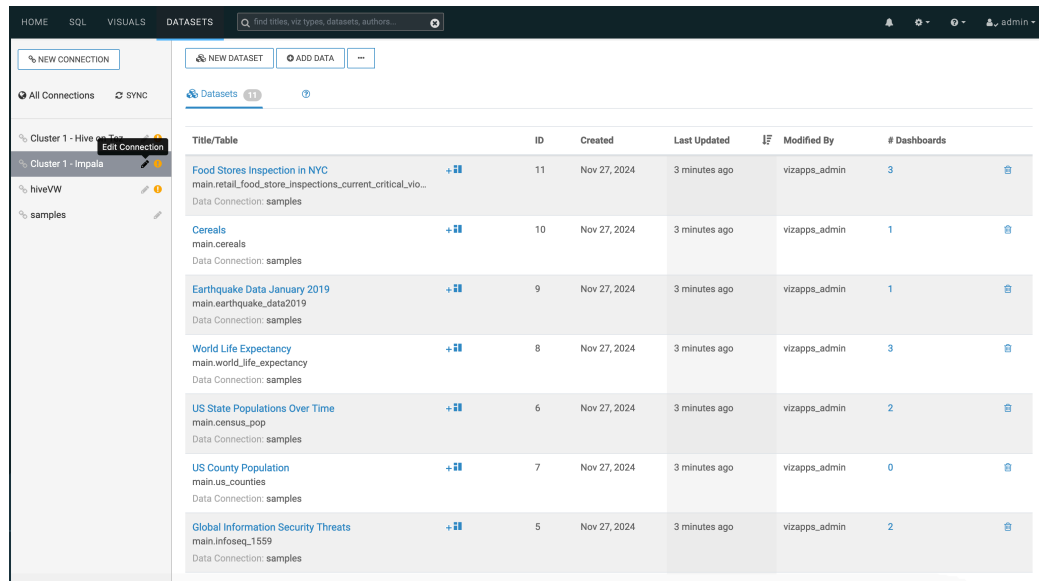


Cloudera Data Visualization on premises for Hive and Impala

For starting Cloudera Data Visualization on premises for Hive and Impala switch to Kerberos Authentication mode in the **Edit Data Connection** page:

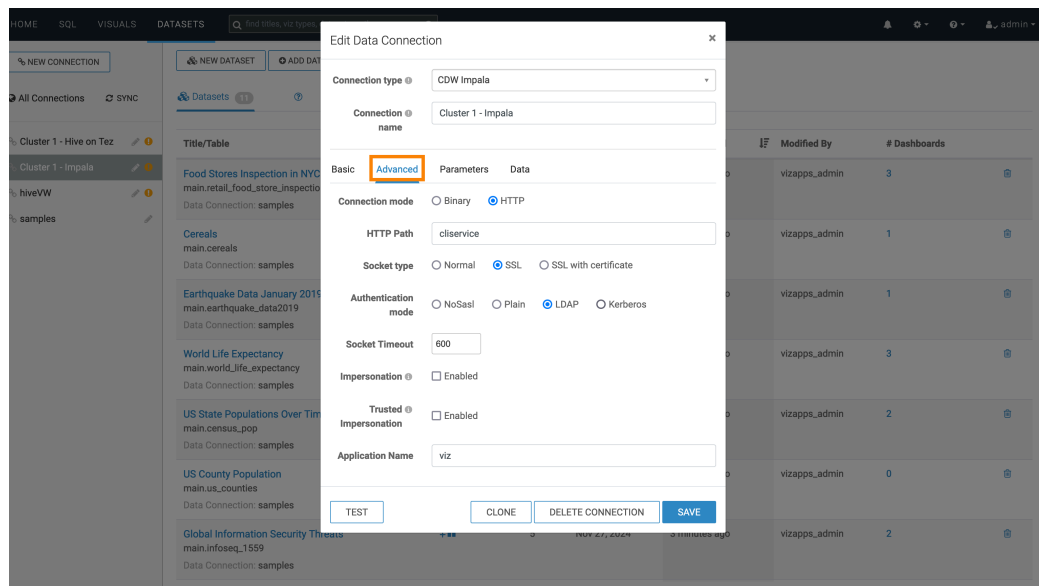
1. Select the required connection in the left navigation and click Edit Connection.

Figure 3: Selecting the required connection for editing



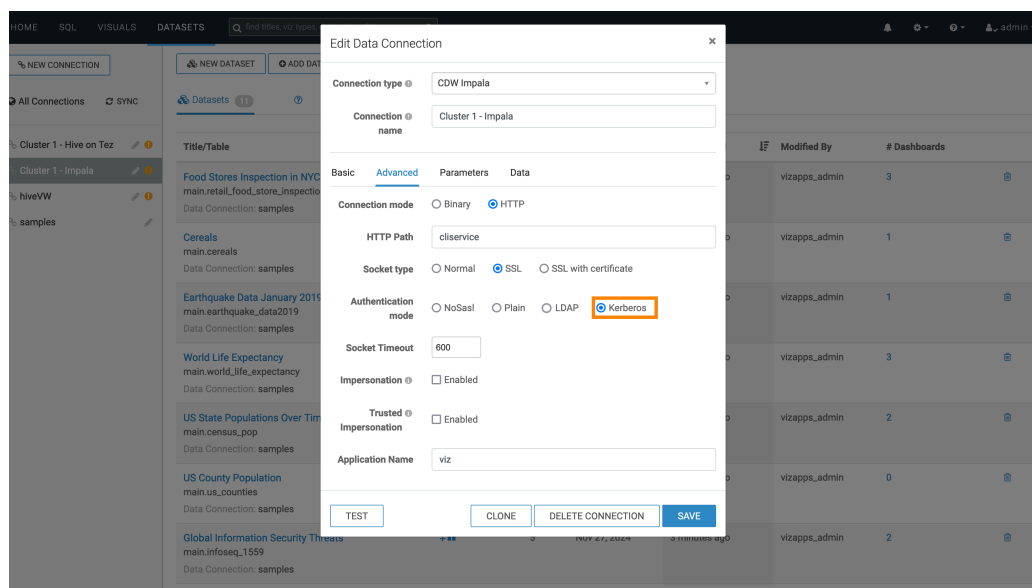
2. Select the **Advanced** tab in the **Edit Data Connection** window.

Figure 4: Selecting the Advanced tab



3. Select Kerberos as Authentication mode in the **Edit Data Connection** window.

Figure 5: Select Kerberos for authentication



Workaround:

For details see, [Workarounds with Cloudera Data Visualization with Hive and Impala](#).

DSE-38164: AMP deployment is not possible in the proxy environment in Red Hat Openshift setup

When the model build is run, the proxy values are referenced from a file path and not from the environment variables. Consequently, the proxy values are not set properly for the models.

Workaround:

A new kubernetes secret needs to be created with the proxy values. Edit the web deployment and add this secret to the path specified.

1. Connect to the kubernetes cluster node and export KUBECONFIG.
2. Copy the script to the node and save it to file as enableProxy.sh

```
#!/bin/bash

# Function to display usage
usage() {
  echo "Usage: $0 -n namespace"
  exit 1
}

# Parse command-line arguments
while getopts ":n:" opt; do
  case ${opt} in
    n )
      NAMESPACE=$OPTARG
      ;;
    \? )
      echo "Invalid option: $OPTARG" 1>&2
      usage
      ;;
    : )
      echo "Invalid option: $OPTARG requires an argument" 1>&2
      usage
      ;;
  esac
done
```

```

    esac
done
shift $((OPTIND -1))

# Check if namespace is provided
if [ -z "$NAMESPACE" ]; then
    usage
fi

CONFIGMAP_NAME="cml-proxy-config"
SECRET_NAME="proxy-config"
NEW_SECRET_NAME="proxy-config-new"
DEPLOYMENT_NAME="web"
DATA_KEY="no_proxy"
SECRET_FIELD="no.proxy"
CONTAINER_NAME="web"
MOUNT_PATH="/config/proxy"

# Step 1: Retrieve the specified data from the ConfigMap
configmap_value=$(kubectl get configmap "$CONFIGMAP_NAME" -n
"$NAMESPACE" -o json | jq -r --arg key "$DATA_KEY" '.data[$
key]')
if [ -z "$configmap_value" ]; then
    echo "Error: Key '$DATA_KEY' not found in ConfigMap '$CONFIG
MAP_NAME'."
    exit 1
fi
echo "Retrieved value from ConfigMap: $configmap_value"
# Step 3: Retrieve the existing Secret
existing_secret=$(kubectl get secret "$SECRET_NAME" -n "$NAM
ESPACE" -o json)
if [ -z "$existing_secret" ]; then
    echo "Error: Secret '$SECRET_NAME' not found."
    exit 1
fi

# Step 4: Modify the specified field in the Secret with the va
lue from the ConfigMap
modified_secret_data=$(echo "$existing_secret" | jq --arg fie
ld "$SECRET_FIELD" --arg value "$configmap_value" '
.data[$field] = ($value | @base64)'
)
# Step 5: Create a new Secret with the modified data
temp_file=$(mktemp)
echo "$modified_secret_data" | jq '.metadata.name = "'$NEW_SEC
RET_NAME'" | del(.metadata.uid, .metadata.resourceVersion, .
metadata.creationTimestamp, .metadata.annotations)' > $temp_
file
kubectl apply -f $temp_file

# Clean up
rm $temp_file

echo "New secret '$NEW_SECRET_NAME' created successfully."

# Step 6: Mount the new Secret to the 'web' container in the
specified Deployment
# Get the current deployment spec
deployment_spec=$(kubectl get deployment "$DEPLOYMENT_NAME" -n
"$NAMESPACE" -o json)
if [ -z "$deployment_spec" ]; then
    echo "Error: Deployment '$DEPLOYMENT_NAME' not found."
    exit 1
fi

```



```

# Extract the existing volumes
existing_volumes=$(echo "$deployment_spec" | jq -r '.spec.template.spec.volumes // []')

# Find the index of the "web" container
container_index=$(echo "$deployment_spec" | jq -r --arg CONTAINER_NAME "$CONTAINER_NAME" '
  .spec.template.spec.containers | map(.name == $CONTAINER_NAME) | index(true) // "null"
')

# Check if the container named "web" exists
if [ "$container_index" == "null" ]; then
  echo "Error: Container with name '$CONTAINER_NAME' not found in deployment '$DEPLOYMENT_NAME'."
  exit 1
fi

# Extract the volumeMounts for the "web" container
existing_volume_mounts=$(echo "$deployment_spec" | jq -r --arg json_index "$container_index" '
  .spec.template.spec.containers[$index].volumeMounts // []
')

# Define the new volume and volumeMount
new_volume=$(echo "[{\"name\": \"$NEW_SECRET_NAME\", \"secret\": {\"secretName\": \"$NEW_SECRET_NAME\"}}]" | jq -r .)
new_volume_mount=$(echo "[{\"name\": \"$NEW_SECRET_NAME\", \"mountPath\": \"$MOUNT_PATH\"}]" | jq -r .)

# Combine existing and new volumes
updated_volumes=$(echo "$existing_volumes $new_volume" | jq -s 'add')
updated_volume_mounts=$(echo "$existing_volume_mounts $new_volume_mount" | jq -s 'add')
# Patch the deployment with the new volumes and volume mounts
kubectl patch deployment "$DEPLOYMENT_NAME" -n "$NAMESPACE" --type='json' -p="
[
  {
    \"op\": \"replace\",
    \"path\": \"/spec/template/spec/volumes\",
    \"value\": $updated_volumes
  },
  {
    \"op\": \"replace\",
    \"path\": \"/spec/template/spec/containers/$container_index/volumeMounts\",
    \"value\": $updated_volume_mounts
  }
]"

echo "New secret '$NEW_SECRET_NAME' mounted to the 'web' container in the deployment '$DEPLOYMENT_NAME' successfully."

```

3. Run the script with the following command and change the Cloudera AI namespace:

```

chmod +x enableProxy.sh && ./enableProxy.sh.sh -n <cml-namespace>

```

4. Monitor the web pods under the Cloudera AI namespace.

OPX-5766: The system is reaching the configured limit of inotify instances

Running too many sessions or nodes might lead to crashing fluentbit sidecar containers. Cloudera AI reports that the session is running but the users might experience limitations, that is, for example, Cloudera AI Workbench does not work. This is because the system reaches the configured limit of inotify instances.

Workaround:

The nodes must have the system variable `fs.inotify.max_user_instances` configured to a reasonable value.

This can be set by running the following in the node:

```
sudo sysctl fs.inotify.max_user_instances=8192
sudo sysctl -p
```

To persist the above configuration with node restarts, add the below line to `/etc/sysctl.conf`:

```
fs.inotify.max_user_instances=8192
```

DSE-42079: Cloudera AI Workbenches are not compatible with the *Unified Timezone* feature

When you enable the *Unified timezone* feature, the Cloudera Embedded Container Service cluster timezone is synchronized with the Cloudera Manager Base time zone, and the Cloudera AI sessions will fail to launch with Exit Code 34. You will also see timestamp discrepancies with workloads. For more information, see [ENGESC-28507](#).

Workaround:

If you use Cloudera AI Workbenches, disable the *Unified Timezone* feature by following the instructions in [Cloudera Embedded Container Service unified time zone](#).

DSE-34314: Templating Cloudera Base AutoTLS properties needed in Cloudera AI sessions

Spark pushdown does not work if Atlas and Spark Lineage checks are enabled for Spark.

Workaround:

To use Spark pushdown, you need to disable Atlas and Spark Lineage checks in the Spark configurations.

DSE-41424: Better handling for data connection validation errors

In Cloudera AI on premises, Cloudera AI Workbench fails to start under the following conditions:

- If the `HIVE_ON_TEZ` service is not present, or if the following configurations are missing:
 - `hiveserver2_load_balancer`
 - `hive.server2.transport.mode`
 - `Kerberos_princ_name`
- If the `IMPALA` service is not present on your Base cluster, or if the following configurations are missing:
 - `Hs2_port` or `hs2_http_port`
 - `Kerberos_princ_name`

Workaround:

1. Disable the data connection auto-discovery feature.

```
kubectl exec -it <db-pod> -n <workspace-namespace>
psql -U sense
# set the flag to false
sense=# update site_config set enable_discover_dataconn_res
tart = false;
UPDATE 1
```

- Restart web pods to apply changes.



Note: The Data connection discovery button may not function properly until the invalid URLs in the database are corrected.

- Check and ensure that there are no incorrect entries in the database.
- Back up the `workspace_data_connections` and the `project_data_connections` tables.

```
[postgres@db-0 /]$ pg_dump -U sense -t workspace_data_connections -f /data-versioned/workspace_data_connections_dump.sql
[postgres@db-0 /]$ pg_dump -U sense -t project_data_connections -f /data-versioned/project_data_connections_dump.sql
```

- Delete the tables.

```
DELETE FROM workspace_data_connections;
DELETE FROM project_data_connections;
```

Fixed issues in 1.5.4 SP1

This section lists issues fixed in this release for Cloudera AI on premises.

DSE-39287: Job configuration update error on parameter type being invalid

The update on the job setting's workflow is fixed in on premises environments without GPUs configured.

DSE-38404: Model Monitoring AMP task failing due to a restart of the model metrics pod

The container in the Model Metrics pod restarted due to insufficient memory, causing a 502 Bad Gateway error on `cdsw.read_metrics(**ipt)`. The memory limit for the container was set to 100 MB, which has now been increased to 1Gi.

Cloudera AI on premises 1.5.4 SP2

Review the features, fixes, and known issues in the Cloudera AI 1.5.4 Service Pack 2 release.

What's new in 1.5.4 SP2

Cloudera on premises 1.5.4 SP2 includes the following features for Cloudera AI.

Introducing auto scaling performance-critical components

Some performance-critical components in Cloudera AI automatically scale based on demand, enhancing the scalability and the reliability of the product. To enable utilising this feature, we have increased the recommended resource requirements for Cloudera AI by 16 cores.

Known issues in 1.5.4 SP2

You might run into some known issues while using Cloudera AI on premises.

DSE-42079: Cloudera AI Workbenches are not compatible with the *Unified timezone* feature

When you enable the *Unified timezone* feature, the Cloudera Embedded Container Service cluster timezone is synchronized with the Cloudera Manager Base time zone, and the Cloudera AI sessions fails to launch with Exit Code 34. Timestamp discrepancies with workloads are displayed.

Workaround:

If you use Cloudera AI Workbenches, disable the *Unified Timezone* feature by following the instructions in [Cloudera Embedded Container Service unified time zone](#).

DSE-41757: Python workloads running multiple-line comments or strings might fail

Python workloads running multiple-line comments or strings might fail to run when using the Workbench Editor.

Workaround:

Run the code using the PBJ Workbench Editor.

DSE-42509: Creating a project does not work on NTP/Airgap proxy when using private repository with SSH key and SSH URL

When creating a project on NTP/Airgap proxy, using a private repository with an SSH key and an SSH URL, an error message is displayed that the project cannot be created.

No workaround is available.

DSE-42510: Fetching Cloudera or Huggingface AMP catalog is failing on NTP/Airgap proxy

When fetching Cloudera or Huggingface Cloudera Accelerators for Machine Learning Project (AMP) catalog on NTP/Airgap proxy, an error message is displayed Error fetching AMP catalog source URL.

No workaround is available.

Fixed issues in 1.5.4 SP2

This section lists issues fixed in this release for Cloudera AI on premises.

DSE-40909: Disabled 'Run experiments' function does not work as expected

When the Administrator disabled 'running experiments' in Cloudera AI Workbench Site Administration Settings, the user was still able to see the **Experiments** menu item in the global- and project-level navigation.

With the help of implementing the sufficient conditions the UI can now hide the disabled function.

DSE-38645: spark.yarn.jars are not set properly in Cloudera AI on premises Spark Pushdown mode

In case of Spark pushdown, the spark.yarn.jars parameter can indicate which jars in the driver shall be transported to YARN executors. Without this parameter set, the YARN executors use the Spark jars from the base cluster provided by Cloudera Distribution of Spark. This can lead to version mismatch between the Spark driver (used Cloudera Data Engineering-provided jars) and the executors.

Now, the spark.yarn.jars parameter is set to use jars from the Cloudera AI driver (Cloudera AI session) to keep the versions matched between the Spark driver and the executors.

DSE-39798: API v1 stop session endpoint shall perform authentication check

Previously, users could stop sessions under projects that they were not authorized to access using the session's Universally unique identifier (UUID). This issue is now resolved.

DSE-41431: Register engine stopped status

When the reconciler is overloaded with a large number of events, the deleted status is still propagated to ensure that engines do not remain in 'Stopping' status.

DSE-41424: Better handling for data connection validation errors

In Cloudera AI on premises 1.5.4 SP1 Cloudera AI Workbench failed to start under the following conditions:

- If the HIVE_ON_TEZ service was not present or the following configurations were missing:
 - hiveserver2_load_balancer
 - hive.server2.transport.mode
 - Kerberos_princ_name

- If the IMPALA service was not present on your Base cluster or the following configurations were missing:
 - Hs2_port or hs2_http_port
 - Kerberos_princ_name

This issue has been fixed.

DSE-41218: Restrict secrets and ingress access from User Service Account role

User access to Kubernetes secrets and ingresses in their own user namespace has been removed.

DSE-34314: Cloudera Base AutoTLS property set incorrectly

The correction to Cloudera Base AutoTLS configurations resulted in the atlas.kafka.ssl.truststore.location value to be set incorrectly.

This issue has been fixed.